

RISK, AMBIGUITY, AND THE SAVAGE AXIOMS:
COMMENT *

HARRY V. ROBERTS

I. ARE THERE UNCERTAINTIES THAT ARE NOT RISKS?

Daniel Ellsberg has recently revived interest in the distinction between "measurable uncertainties," or "risks," and "unmeasurable uncertainties," or "uncertainties."¹ Ellsberg argues that vagueness about probabilities can lead people to violate the axioms of consistent behavior, the "Savage axioms," upon which modern Bayesian decision theory and inference is based. He contends that some of these violations are conscious and deliberate, not careless mistakes that would be retracted after careful reflection. From this Ellsberg concludes that Bayesian theory is applicable to decision-making under risk but not to decision-making under uncertainty. For dealing with uncertainty a more elaborate theory is needed, and Ellsberg proposes such a theory.

The problem of uncertainty that concerns Ellsberg is real and important, but his diagnosis and therapy seem to me to be a step in the wrong direction. People who are reluctant to use Bayesian theory are likely to be even more reluctant to use Ellsberg's modification of it. The modification demands additional formal assessments of things that are hard to assess, and difficulties of formal assessments are really at the heart of almost all objections to the Bayesian approach, at least by skeptical statisticians.² I shall not, however, press my objections to Ellsberg's therapy. Rather, I shall argue that his diagnosis is wrong, and that the "risk-uncertainty" dichotomy is not really a fruitful one from a normative point of view.

* Valuable comments on earlier drafts of this paper have been given by Gordon Antelman, Selwyn Becker, Ward Edwards, Bruno de Finetti, Giora Hanach, William Kruskal, John Pratt, L. J. Savage, Robert Schlaifer, Lester Telser, and members of the Econometrics Workshop, University of Chicago.

1. Daniel Ellsberg, "Risk, Ambiguity, and the Savage Axioms," this *Journal*, LXXV (Nov. 1961), 643-69. See also William Fellner, "Distortion of Subjective Probabilities as a Reaction to Uncertainty," this *Journal*, LXXV (Nov. 1961), 670-89.

2. The other objection commonly advanced by skeptical statisticians is that Bayesian methods, applied to some particular problem, might involve prohibitive computation or prohibitive analytical difficulty. This objection seems less fundamental than the one mentioned in the text. If the only difficulties were computational or analytical, the Bayesian framework would still seem useful in suggesting approximate solutions or in evaluating rules-of-thumb.

II. VAGUENESS

People often feel that their judgments about probabilities are too vague, ambiguous, volatile, uneasy, or shaky to be taken seriously as a basis for a formal analysis. While these adjectives perhaps have different shades of meaning, let us fix on "vague" as a basis for further discussion. Savage himself has long recognized the problem occasioned by "vagueness." In Savage's terminology, a person is vague about the probability assigned to a single trial if he cannot obtain from himself a clear answer as to what probability to assign to it. He is vague about a probability distribution if introspection fails to reveal clearly what the distribution is. Savage comments, "Some people see the vagueness phenomenon as an objection; I see it as a truth, sometimes unpleasant but not curable by a new theory."³

Vagueness in probability judgments is a matter of degree; we are always more or less vague. The important question is whether or not we are so vague about a probability as to be unwilling to make an assessment of it. If we are willing to make an assessment, we mean that we are willing to incorporate the assessment into a formal decision-theoretic analysis. We can always write down a number, but we may or may not feel that the number is worth basing anything on. This is tantamount to saying that we may wish to make a formal analysis or that we may wish to fall back on an informal one.⁴ It is the distinction between formal and informal analysis, rather than between risk and uncertainty, that seems really helpful.

Vagueness should not be confused with the difficulty or unpleasantness of making assessments. Decisions, however arrived at, can be tough and unpleasant. We may have afterthoughts and misgivings when we make them. But some decisions are unavoidable and must be made somehow. If we do not face up to them, we avoid unpleasantness at the expense of a retreat from our problems. Probability assessments, as one ingredient in decisions, can also be unpleasant. They may sometimes be bypassed or partially bypassed, but the decision to do so should not be based on unpleasantness, except insofar as unpleasantness can be considered a cost to be taken into account in making the decision.

3. L. J. Savage, "Bayesian Statistics," to appear in *Decision and Information Processes* (New York: Macmillan, forthcoming).

4. Of course, even an abortive attempt at formal analysis may be helpful in a subsequent informal analysis by clarifying the issues that need to be mulled over, or by focusing suspicion on any snap decision that might casually emerge from informal analysis alone.

The problem of vagueness is most serious in complicated probability assessments, such as an assessment of a multivariate prior distribution in an analysis of variance problem. The essential issues, however, can be illustrated by the probabilities associated with a Bernoulli process. We therefore present the following terminology. Denote the probability of success on a single Bernoulli trial by p . If we think of p as a random variable that has a distribution of its own, we write \tilde{p} (p-tilde). If we somehow assess the distribution of \tilde{p} , the assessment of p is the expected value of this distribution, $E(\tilde{p})$.⁵ We could be vague about the distribution of \tilde{p} yet not about $E(\tilde{p})$. If we were vague about $E(\tilde{p})$, we would necessarily be vague about the distribution of \tilde{p} .

It is tempting to say that vagueness has something to do with the dispersion of the distribution of \tilde{p} . While we may find it more difficult or unpleasant to make a judgment about a dispersed distribution, we are not vague if we do make a judgment. A dispersed distribution of \tilde{p} will, however, be more sensitive to revision in the face of sample data from the Bernoulli process.

III. ELLSBERG'S PARADOX

Ellsberg's first example contains all the ingredients needed for further discussion, so we quote his description in full.

Consider the following experiment. Let us suppose that you confront two urns containing red and black balls, from one of which a ball will be drawn at random. To "bet on Red₁" will mean that you choose to draw from Urn I; and that you will receive a prize a (say, \$100) if you draw a red ball ("if Red₁ occurs") and a smaller amount b (say, \$0) if you draw a black ("if not-Red₁ occurs").

You have the following information. Urn I contains 100 red and black balls, but in a ratio entirely unknown to you; there may be from 0 to 100 red balls. In Urn II, you confirm that there are exactly 50 red and 50 black balls. An observer—who, let us say, is ignorant of the state of your information about the urns—sets out to measure your subjective probabilities by interrogating you as to your preferences in the following pairs of gambles:

1. "Which do you prefer to bet on, Red₁ or Black₁: or are you indifferent?"

5. Since there seems to be some confusion on this point, a proof may be in order. We use the discrete case to show the method:

$$\begin{aligned}
 P(\text{success on one trial} \mid \text{distribution of } \tilde{p}) &= \sum_p P(\tilde{p} = p) P(\text{success} \mid p) \\
 &= \sum_p P(\tilde{p} = p) \cdot p \\
 &= E(\tilde{p})
 \end{aligned}$$

ent?" That is, drawing a ball from Urn I, on which "event" do you prefer the \$100 stake, red or black: or do you care?

2. "Which would you prefer to bet on, Red_{II} or Black_{II}?"
3. "Which do you prefer to bet on, Red_I or Red_{II}?"
4. "Which do you prefer to bet on, Black_I or Black_{II}?"⁶

Ellsberg reports that many people are indifferent in questions 1 and 2, yet express a preference for Red_{II} in question 3 and Black_{II} in question 4. (Presumably this pattern of responses reflects a preference for dealing with a "known" distribution.) He goes on to demonstrate that this set of answers cannot be reconciled with the Savage axioms, *given Ellsberg's interpretation of the problem that he has formulated*. A person must either modify his answers or violate the axioms. Ellsberg argues essentially that the intuitive conviction underlying the set of answers may be stronger than the intuitive conviction of the axioms, in which case the axioms should yield. To me the intuitive conviction of the axioms is strong and direct. But my experience in trying to apply them has convinced me of two things: (1) it is easy to make mistakes, or at least to misinterpret the problem to be solved; (2) I have thus far succeeded in tracing all my violations of the axioms to mistakes or misinterpretations. My recommendation to a person giving the answers reported by Ellsberg is this: "Put yourself temporarily within the Bayesian framework and see what a careful Bayesian analysis can contribute to your understanding of your answers."

An analogy may be helpful. Optical illusions are well known to distort perception, and the study of such illusions has been an important and fruitful area of psychology. We all know that there are circumstances in which we cannot trust our own eyes. If I watch a magician, the plain evidence in front of me tends to undermine my confidence in certain natural laws. Yet I firmly believe that, whether or not I can find it, there *is* a natural explanation. It does not occur to me to discard or modify well-known laws of nature.

The analogy is unfair in that it seems to suggest that we should put the same trust in the Savage axioms as we would, say, in the law of conservation of mass and energy.⁷ But the analogy is useful

6. *Op. cit.*, pp. 650-51.

7. Ellsberg points out the danger that the Savage axioms may somehow "abstract away from vital considerations." *Op. cit.*, p. 669. Luce and Raiffa remind us, ". . . before committing ourselves on the acceptability of a set of axioms . . . , we are well advised to investigate some of the consequences with an idea of ferreting out the hidden jokers." R. Duncan Luce and Howard Raiffa, *Games and Decisions* (New York: Wiley, 1957), p. 124. My own confidence in the Savage axioms stems mainly from their fruitfulness in statistical problems, although their direct intuitive appeal is also important to me.

in suggesting that apparent violations of the Savage axioms might be fruitfully studied in the same spirit that we might study the moon illusion, the fact that the moon looks larger very near the horizon than when it is higher in the sky, a phenomenon that has invited explanation for centuries.⁸

We begin by noting that each of the four answers mentioned by Ellsberg could be interpreted as a probability assessment, and that the probability assessments are *not* vague. The trouble is that the set of answers displays inconsistent probability assessments. Thus:

1. $P(\text{Red}_I) = P(\text{Black}_I) = .5$
2. $P(\text{Red}_{II}) = P(\text{Black}_{II}) = .5$
3. $P(\text{Red}_{II}) > P(\text{Red}_I)$
4. $P(\text{Black}_{II}) > P(\text{Black}_I)$.

On any one question, the subject *has* obtained from himself a clear answer about the probability to be assigned to an individual trial. He is presumably prepared to act upon it in making a decision.

Two important lessons can be drawn from this. (1) Suppose that a decision *must* be made about an individual trial, and that the decision turns only on how a probability would be assessed if he tried to assess it. Then to say that the subject is vague is to say that his power of decision is paralyzed. He can afford the luxury of vagueness only if the decision can somehow be deferred or avoided. (2) By noticing his inconsistent probability assessments, he can study the inconsistencies carefully to see which assessments he holds to and which ones should yield. This process lies wholly within, and indeed typifies, the Bayesian approach, for which internal consistency is the major part of "rationality."

In assessing the probability p of "success" on an individual trial one approach is to assess the distribution of \tilde{p} and compute $E(\tilde{p})$ from this distribution. A second approach is to assess $E(\tilde{p})$ directly. There is an option here, and there is no *logical* principle to say which approach should be used. (Indeed the subject might try it *both* ways, and, if he discovers inconsistencies, reflect further to learn better his own mind.) Let $p_I = P(\text{Red}_I)$, $q_I = 1 - p_I$, $p_{II} = P(\text{Red}_{II})$, and $q_{II} = 1 - p_{II}$. Presumably everyone agrees, without any problem of vagueness, that $E(\tilde{p}_{II}) = .5$ and that for all practical purposes the distribution of \tilde{p}_{II} concentrates mass unity at $p_{II} = .5$. The answer given by the subject to question 1 suggests that, for him, $E(\tilde{p}_I) = .5$. The distribution of \tilde{p}_I is pre-

8. Lloyd Kaufman and Irvin Rock, "The Moon Illusion," *Scientific American*, Vol. 207 (No. 1, 1962), pp. 120-30.

sumably *not* concentrated at .5 but rather spread out in some way that leads to an expectation of .5. The subject could quite properly have claimed vagueness about all other details of this distribution, since it was not necessary to remove vagueness about the distribution in order to assess its expected value.

In answering questions 3 and 4, the subject, who, recall, is now trying to reapproach the problem in the spirit of the Savage axioms, finds that he has assessed

$$E(\tilde{p}_{II}) > E(\tilde{p}_I)$$

and

$$E(\tilde{q}_{II}) > E(\tilde{q}_I).$$

It is hard to see that questions 3 and 4 have given him any insight about the urns that he did not have when he answered 1 and 2. Yet his intuitive feelings about his answers to 3 and 4 are still strong and he is tempted, with Ellsberg, to feel that the Savage axioms will have to yield. But before giving up on the axioms a more careful examination of the answers to 3 and 4 is wise.

In preparation for this, let me state my own understanding of certain points about Ellsberg's problem that are not made explicit by his wording. I assume that Ellsberg meant that the problem should be analyzed as a single, unique choice; that the utility of the outcomes of the games depended solely on money; that the money outcomes depended only on the outcomes of an individual trial; and that there are no concealed possibilities, such as the possibility of buying sample information. With this in mind let us consider the various ways of resolving the preference for II over I in spite of indifference as to color within I or II separately.

(1) Ellsberg's analysis of his paradox is based on the assumption that the utilities of outcomes are a function only of the monetary consequences. But this is not necessarily the assumption that the subject was making. As just one illustration,⁹ the subject might feel that his choice of Red_I could lead to unpleasant second guesses by someone who observed the experiment: he could be criticized, however unfairly, for not taking an apparently "safe" course of action (Red_{II}) if he lost by taking an "unsafe" one (Red_I). An analysis of utility, formal or informal, might reveal the reason for the subject's willingness to pay more for Red_{II} than Red_I, and similarly for Black_{II} over Black_I. At the same time he could resolve his confusion about probabilities by deciding that he really feels

9. For another in the same vein, the subject may fear that Urn I might contribute to an ulcer, regardless of what a rational analysis of other aspects of the problem may suggest.

$P(\text{Red}_I) = P(\text{Red}_{II}) = P(\text{Black}_I) = P(\text{Black}_{II}) = .5$, without taint of vagueness.

(2) The subject may have assumed erroneously that the monetary outcomes are a function of p rather than of the outcome of a single trial. For example, he may subconsciously have been thinking that the payoff was proportional to the number of red balls in the chosen urn. If each red ball was worth \$1, the choice Red_{II} is worth \$50 *for sure* in this game. The choice Red_I has, by the answer to question 1, an expected *monetary* value of \$50. But even without a careful assessment of the distribution of \tilde{p}_I , it will be clear that expected *utility* will be less (assuming a conditional utility curve representing "risk aversion") than the utility of \$50 certain. For this game, therefore, the choice Red_{II} would be indicated by a utility argument and no inconsistent probability assessments need be involved. It does not to me seem fantastic that the subject should get mixed up about which game he is playing. I have done the experiment in statistics classes and encountered the following kind of argument. "If I pick Red_I , I would be completely out of luck if there were no red balls in the urn." Students seem less ready to see the fact that one is equally out of luck if he picks Red_{II} and a black ball is drawn.

Closely related logically to the previous paragraph, but more subtle—the subject may not want to analyze the problem as a single unique choice. He may have it in mind that he will play the game again. One simple example is easily analyzed and instructive. Suppose that the subject must decide now on his choice of red or black for 100 independent repetitions either from Urn I or Urn II. By independent, I mean in the sense of drawing *with* replacement repeated samples of size 1. It is now quite possible that he is indifferent between Red and Black for either urn separately, but prefers Urn II to Urn I. This is because: (a) Monetary return is a function of 100 trials rather than one trial. (b) Given a uniform prior¹ for \tilde{p}_I , the distribution of monetary return has the same expected value for each urn but approximately three times the standard deviation for Urn I. (c) In the presence of risk aversion, larger dispersion of the distribution of monetary return implies smaller expected utility.

(3) The subject may now notice that the "cost of uncertainty," in the technical sense, is higher for Urn I than Urn II. That is, if he had the option of paying for sample information before making a

1. Introduced only for concreteness. Any nondegenerate prior leads to the same qualitative conclusion.

choice, he would pay nothing at all for sample information from Urn II, while he might well pay something for sample information from Urn I. If we assume that the subject is not vague about the distribution of \tilde{p}_I and is willing to assess it for rough calculation as a continuous uniform distribution² from 0 to 1, and that he is guided by expected monetary value, it turns out that he should be willing to pay up to \$16.67 for a sample of *one* ball before making his choice. For, if the sample turns out red, Bayes' theorem will cause him to revise his expectation for \tilde{p}_I from .5000 to .6667. If it turns out black, he will revise from .5000 to .3333. In either case, the expected income for the choice that looks best after the sample (red if red, black if black) is increased from \$50 to \$66.67. The question of how much he would be willing to pay for sample information is logically irrelevant since none of Ellsberg's questions admit this option, but it is easy to be confused.³

(4) The subject may decide that the choice of Red_{II} over Red_I would be held only if all other considerations were evenly balanced, but that upon reflection he would be willing to pay nothing at all for the privilege to choose Red_{II} . His original preference was really a second order preference, so to speak, and should not lead him to abandon the axioms.

(5) The subject may have had a much harder time in answering question 1 than question 2. The probability assessment of question 2 was trivial, whereas the assessment of $E(\tilde{p}_I)$ in question 1 was difficult. Granted that he cannot afford the luxury of vagueness under the circumstances, the subject would certainly not have an easy mind about what he had done. He might feel that if he had to make the assessment repeatedly, he would be wildly volatile in his choices according to his latest feelings about the whims of the experimenter who had filled up Urn I. It is easy to carry over this feeling into questions 3 and 4, and want to have no further part of Urn 1.

(6) The subject may realize he is vague, in Savage's sense, about the *distribution* of \tilde{p}_I , and that he failed to realize that this distribution is irrelevant to either question 3 or 4.

I cannot claim dogmatically that all subjects are confused by one or more of the six points just suggested. Presumably psychological experiments could be designed to shed further light on the

2. See preceding footnote.

3. In commenting on this paragraph, Robert Schlaifer described a theoretical problem for which he and Arthur Schleifer, Jr., initially made a false conjecture about the solution because they were confused by the existence of a cost of uncertainty that was not relevant to the problem they were investigating.

question. From the point of view of psychology, errors of perception are interesting and knowledge about the causes of these errors is valuable. My interest, however, is mainly normative. I contend that it is irresponsible and unwise for a subject to persist in his four original answers unless he has first carefully examined these six points and if he has not looked for other possible explanations in the same vein. (My list is intended to be suggestive rather than exhaustive.)⁴ Almost everyone seems to accept the Savage axioms in the absence of vagueness about probability assessments. Violation of the axioms because of vagueness should not be contemplated lightheartedly if the decisions or inferences involved are taken seriously.

IV. DISCUSSION AND CONCLUSION

It is hard to see any important role for vagueness in Ellsberg's paradox, at least for a person who makes definite choices for all four questions. Yet vagueness and the threat of it are often important problems in application of Bayesian decision theory. Irresponsible handling of vagueness renders the whole apparatus hazardous, just as would any form of self-deception in any scheme of decision-making. It would be helpful to look more closely at the process of arriving at assessments of probabilities to see just how vagueness may enter in and what can be done about it. But this would lead to a rather technical chapter in Bayesian statistics, one that can be written convincingly in the present state of the art only for relatively simple problems. While I shall not here undertake even a summary of this chapter, it is worth pointing out my conviction that people skeptical of the Bayesian approach are likely to be convinced only by seeing concrete Bayesian methods that can cope with serious problems. A Bayesian can only wish that skeptics would not make firm judgments until they have given the approach a fair chance.

4. Ward Edwards, after reading an earlier draft of this paper, did some informal experimentation. He comments: "I found the opinion that a seventh source of error, different from the six you list, was controlling. That is, a simple preference for more information . . . over less. Such a preference is a natural consequence of the experimental fact that information cannot hurt and almost always helps in decision-making." Gordon Antelman suggested in the same vein: "Subject may be pretty sure he doesn't know how to analyze the game and yet be fairly certain that Urn II *is at least* as good as Urn I."

A different line of reasoning, based on mixed strategies, also tends to reinforce the appeal of the Savage axioms in examples closely related to, though not identical with, Ellsberg's. Raiffa points out that the subject can assure himself an objective probability of .5 for Urn I by tossing a fair coin to decide which color to pick. See Howard Raiffa, "Risk, Ambiguity, and the Savage Axioms: Comment," this *Journal*, LXXV (Nov. 1961), 690-94.

Ellsberg's paradox serves as a useful example for showing that one's first intuitive judgments may be analyzed and illuminated by the very approach that they seem at first to threaten. I do not contend that all subjects would hold to the Savage axioms as a result of such analysis, but I believe, partly on the basis of informal experimentation, that many of them would.

In any event the remedy for vagueness is an honest attempt to recognize genuine vagueness, to deal with it directly if possible, or to bypass it skillfully by less formal and complete analyses. My personal opinion is that the problem of vagueness will be most successfully met in situations in which at least part of the information comes from sample data, that is, numbers generated by such processes as the Bernoulli, Poisson, or normal. When sample data are absent or when vagueness threatens our attempts to assess the sampling process, the role of formal analysis may have great conceptual value, for example, in disentangling the *probabilities* of events from the *utilities* of the consequences of events. The formal approach may hint at good informal ones, as when we graph data in ways that cast light on such assumptions as independence or normality. But the literal application of formal methods is likely to be much more restricted. Even so, the normative value of Bayesian decision theory can be great. If we cannot always eliminate vagueness about the answers, we never need be vague about the right questions to ask.

UNIVERSITY OF CHICAGO

REPLY

DANIEL ELLSBERG

There is so much in these matters on which Professor Roberts and I agree that, if a few summary sentences were slightly different, I would be tempted simply to thank him for supplementing my earlier discussion and underlining certain aspects of it. But that would be rather tactlessly to ignore Roberts' own interpretation of his remarks, which seems quite otherwise. If I am to take issue with him, it must be not with his treatment of my position — as criticism, his comments are admirably restrained, fair-minded and cautious — but with his understanding of the thrust of his own specific comments.

A major point on which there seems no disagreement is the fact and the general pattern of violations of the Savage axioms in connection with my hypothetical choice problems. Of course, in testing the acceptability of normative postulates, hasty, undeliberated choices are in no way conclusive and violation must be considered tentative¹ pending thorough analysis and reflection in the light of all implications of the theory.² "I do not contend," Roberts concludes, "that all subjects would hold to the Savage axioms as a result of such analysis, but I believe, partly on the basis of informal experimentation, that many of them would." That conclusion could summarize my own observations, though I might reverse the emphasis. As I reported in my article,³ responses do vary. Of those who tentatively violate the postulates, some conclude upon reflection or interrogation by unsympathetic critics (sometimes, themselves) that their initial choices were "mistakes" which they now wish to modify; others do not. For further reference, I shall call the former group "transient" violators and the latter, "deliberate" violators.⁴ Roberts' introspective experience with his own violations places him, so far, in the first category; I remain, so far, in the second. But though our own reactions fail to match, I gather that our observations of response varieties have been similar. Our notions of the relative frequency of "deliberate" violation may differ somewhat, and Roberts is more concerned to lower it, if possible. But he does not claim that the violations of his subjects have all so far, like his own, proved transient (I presume he would not fail to mention it if they had), nor does he conjecture that with adequate reflection by all concerned the class of deliberate violators will prove empty. It is only this residual group of deliberate violators who raise prob-

1. Likewise conformity, where the appropriateness of the theory is in question.

2. See "Risk, Ambiguity, and the Savage Axioms," this *Journal*, LXXV (Nov. 1961), 655-56, for comments and admonitions paralleling Roberts'. For a much fuller discussion of the nature, validation and functions of normative theories of choice, see "Risk, Ambiguity and Decision," The RAND Corporation, RM-3543 (forthcoming); again, I think we are in complete agreement on these questions.

3. *Op. cit.*

4. This distinction can be made hollow, to whatever end that serves, by taking the period of "adequate" reflection to be sufficiently short or sufficiently long. For practical purposes, those violators who conclude that they wish to persist in their violating choices after conscientious consideration of such critiques as those of Roberts and of H. Raiffa ("Risk, Ambiguity, and the Savage Axioms: Comment," this *Journal*, LXXV, Nov. 1961, 690-94) may be classed as "deliberate." Although such persons may always change their minds eventually, the Savage axioms do not constitute an appropriate, usable model of their actual, deliberated preferences *now*; if they are to benefit from a normative logic of choice, it must be somewhat different.

lems of normative theory, and my "diagnosis" (that "ambiguity" was a major contributing factor) concerned them exclusively.

Roberts defines "vagueness" in terms that correspond to my notions of "ambiguity" well enough for this particular discussion:

A person is vague about the probability assigned to a single trial if he cannot obtain from himself a clear answer as to what probability to assign to it. He is vague about a probability distribution if introspection fails to reveal clearly what the distribution is . . . we are always more or less vague.⁵

I presume that Roberts agrees that my hypothetical examples tend to induce considerable vagueness of opinion concerning certain alternatives (his subjects would tell him, if his own introspection did not).⁶ Moreover, Roberts emphasizes that the problem of vagueness is "real and important." What he denies is that the evident vagueness of opinion in these instances contributes in any important way to the violations of the Savage postulates that admittedly do occur.⁷

5. Though Roberts furnishes no formal model for this state of mind, his definition immediately suggests one: the whole *set* of probability distributions that are not *ruled out* by those probability comparisons of which the individual feels relatively "sure." This assumes that introspection does reveal a relatively clear answer with respect to *some* probability comparisons; typically, these judgments are in the form of inequalities, statements that the probability of one event is greater (or, not less) than that of another. To say that "we are always more or less vague" may be interpreted to mean that the set of probability distributions compatible with (not excluded by) all our definite probability judgments at a given moment typically contains more than one member. In general, this model would associate a particular event not with a precise probability but with an *interval* of probability-numbers circumscribed by inequalities, and the relation "not less probable than" is regarded as providing only a *partial ordering* among events.

The distribution-set is denoted Y° in "Risk, Ambiguity, and the Savage Axioms" (p. 661). At the time of writing that article, I was unaware that this model of opinion, encompassing "vague" opinion and "ambiguity," had been elaborately developed in earlier, important works by B. O. Koopman ("The Axioms and Algebra of Intuitive Probability," *Annals of Mathematics*, Series 2, Vol. 41, 1940, pp. 269-92) and I. J. Good ("Rational Decisions," *Journal of the Royal Statistical Society*, Series B, Vol. 14, 1952, pp. 107-14; more recently, "Subjective Probability as the Measure of a Non-Measurable Set," *Proceedings of the International Congress for Logic, Methodology and Philosophy of Science*, Stanford University Press, 1962, pp. 319-29). These works are further discussed in "Risk, Ambiguity and Decision."

6. In the specific example he discusses, the subjects are given no explicit information on the ratio of red to black balls in Urn I; in practice, subjects readily report great vagueness of opinion concerning their prospects of winning bets on the color of a ball to be drawn from Urn I (bets on Red; or Black), in contrast to their precise opinions on the probabilities of winning corresponding bets on Urn II, in which the proportion of red to black is known to be 50:50. In terms of the model in the preceding footnote, the opinion of such a subject concerning a drawing from Urn I (or, concerning the *ratio* of red to black in Urn I: an uncertain fact on which the prize does not depend directly but which underlies the subject's evaluation of the bet) must be represented by a large set of distributions, while his opinion concerning a drawing from Urn II may be represented by a single, precise distribution.

7. Why, then, is it "important"? How does vagueness affect decision? I

He concludes: "It is hard to see any important role for vagueness in Ellsberg's paradox, at least for a person who makes definite choices for all four questions." The challenge seems clear enough. But on a closer look, the conflict blurs.

The heart of Roberts' specific critique is a list of considerations that could induce a subject to make choices in the "two-urn example" cited above in a pattern that evidently violates the Savage postulates: specifically, Postulate 2, the Sure-thing Principle. These considerations include possible mistaken beliefs as to the explicit conditions of the bets: for example, the facts that the game is to be played only once, that a single drawing is to be made from the selected urn, and that the money payoff depends only on the color of the ball drawn, not on the proportion of colors in the urn.

Roberts' discussion both before and after this list is presented tends to suggest that *all* the considerations to be mentioned are of the character of the three above: "mistakes, misinterpretations, misconceptions" that may "confuse" a subject but which, once brought into consciousness and made explicit, will not induce him to persist in violations when the problem is properly understood. Moreover, his conclusion implies that vagueness plays no part in these considerations.

But neither of these characterizations of the proposed rationales seems accurate; and even if they were, Roberts' argument would not sustain his conclusion.

The first characterization, if valid, would imply that he has simply compiled a list of various sources of *transient* violations: an believe that Roberts would answer: It affects the difficulty of decision-making, the time and effort required, the pleasantness of the task and one's confidence or uneasiness in the results, the frequency of random errors and "brief" transient violations; but *not*, given enough time for reflection and analysis, the answers one will ultimately choose, hence *not* the acceptability of Savage's normative postulates. In short, vagueness of opinions affects "feelings" more than decisions, promotes indecisiveness and vacillation, and affects decision-making in the same sorts of ways as does *complexity*.

My own view is that in addition to these effects, and partly because of them, vagueness of opinion can affect the choices that seem preferable on thorough reflection, and for some people its influence does lead to deliberate conflicts with the Savage postulates. Vagueness need not be synonymous with indecision, as Roberts implies; *some choices are easy, just because certain opinions are vague* (for many people this is true for the choice between a bet on Red_I and a bet on Red_{II}, i.e., for questions 3 and 4 cited by Roberts). Nor must such people turn to informal analysis, if a more appropriate formal theory can be made available. I think it can.

I proposed *one* candidate in my article and discuss several others, including some I would favor (omitted, regrettably, from the article for reasons of space) in "Risk, Ambiguity and Decision." Most of this "therapy" was antedated by I. J. Good in his remarkable paper, "Rational Decisions" (cited above). For a recent parallel, see C. A. B. Smith, "Consistency in Statistical Inference and Decision," *Journal of the Royal Statistical Society, Series B*, Vol. 23 (1961), pp. 1-25.

interesting and useful effort, but not addressed to *my* problem. Used in self-interrogation, such a list (in which vagueness would surely be but one factor among many) would be helpful in sifting out the transient from the deliberate violators more efficiently. But unless one conjectured that such a sieve would, in practice, show the latter, residual class to be null—and Roberts does not press his argument this far—it would seem to provide, at best, imperfect insight into the “paradox” (in Roberts’ eyes) of persistent, deliberate violation.

What *are* the considerations, in Roberts’ view, that influence the deliberate violator, he whose transgressions of the Savage postulates must be adjudged neither lighthearted, irresponsible, nor unwise? Roberts does not dispute his existence, yet he ventures no explanation; nor does he directly attack the one I propose. In fact, he seems to be silent on the matter. But this is the precise point at issue, if my “diagnosis”—which concerned *only* these residual subjects—is in question!

The assertion that the diagnosis of vagueness has been shown to be “wrong” or irrelevant appears even more puzzling when one considers the actual substance of the points Roberts raises in his critique. In the six numbered passages and two of the footnotes (footnote 9 on p. 333 and footnote 4 on p. 336) toward the end of Section III, I count nine distinct “resolutions” of the pattern of violations in question. In *six* of these, implicitly or explicitly, *vagueness seems to play a criticale role!* I shall comment on the more important of these (the reader is referred to Roberts’ text for the exact import of the propositions cited):

(a) “The subject may have had a much harder time in answering question 1 than question 2.” Really? Why? except that he finds his opinions more *vague*, with respect to question 1! (Question 1 concerns a drawing from Urn I, for which the ratio of red to black has not been specified.) Why else would he “certainly not have an easy mind” about any given resolution of the question; why else might he expect himself to be “wildly volatile in his choices” if he had to make the assessment repeatedly? (Either of these expectations could serve as a fairly adequate working definition of “vagueness” of opinion; but here and below, I rely on Roberts’ definition cited earlier.) As Roberts suggests, “it is easy to carry over this feeling” into a preference for bets based on Urn II (in which the ratio is known to be precisely 50:50). But this is a response to relative *vagueness!* And by the way . . . in what sense does it reflect “mistakes,” confusions, misinterpretation of the prob-

lem? Can, or should, an application of the Savage postulates abstract from such "difficulties" that issue from the very quality of the uncertainties in question?

(b) "The subject may realize he is vague, in Savage's sense, about the distribution" of the *ratio* of Red to Black in Urn I, and fail "to realize that this distribution is irrelevant to either question 3 or 4." But vagueness about the distribution is not conducive to precision or confidence about its mathematical expectation, which is relevant to questions 3 and 4. Roberts mentions earlier that "the answer given by the subject to question 1 suggests that, for him," this expectation = .5; but the *other* answers postulated for him, to questions 3 and 4, are *inconsistent* with this interpretation, as they are with the inference of a single, definite probability distribution over the possible ratios. As Roberts says, the subject "could" claim a precise opinion on the expectation despite vagueness as to the distribution (somewhat implausibly, unless he avows compelling intuitions of symmetry that are not obviously appropriate here); but he need not do so, merely to justify his answer to question 1.⁸ In any case, the role of vagueness is explicit here, whether or not a "mistake" is involved.

(c) *Why* might the subject "feel that his choice of Red_I could lead to unpleasant second guesses by someone who observed the experiment"? Why might he be criticized "for not taking an apparently 'safe' course of action (Red_{II}); is he lost by taking an 'unsafe' one (Red_I)"; on what grounds can bets on Red_I and Red_{II} be discriminated by a potential critic save relative vagueness of accompanying opinions? Indeed, on what other basis can the terms "safe" and "unsafe," as applied here, be interpreted meaningfully? (The very fact that these terms do seem apt in this context deserves some serious thought from Roberts; these examples were partly constructed just to elicit such notions without offering a basis for them in terms of the range, minimum payoff, variance or expected value of a specified distribution.) Even if, as Roberts postulates, the subject's own opinions happen to be precise, it is vagueness — in this case, the anticipated perception and evaluation of vagueness by others — that determines his hypothesized response.

(d) "In the same vein, the subject may fear that Urn I might contribute to an ulcer." Quite: but why Urn I? Obviously, because he sees the vagueness that we see, and that he might expect others

8. The set of distributions representing his vague opinions concerning Urn I may have certain symmetry features (by containing matched pairs of asymmetric distributions) that account for his indifference between bets on Red_I and Black_I; no precise "expectation" need attach to the set.

to see. Even if we abstracted, experimentally, from the problem of anticipated second guesses by others by keeping his choices private, we could not protect him from no less unpleasant second-guessing by himself; from self-reproaches, from *regret* (in a familiar, not a technical, sense that would repay analysis) evoked by losses on Urn I, for reasons intimately associated with vagueness.

Once again, are most of these considerations based on mistakes, confusions, misconceptions? Is the subject wrong to expect the epithet "unsafe" to be attached to Red_I and not to Red_{II} ; and is he foolish to take that into account? Is the ulcer-prone individual mistaken, or arbitrary, to see Urn I as the more threatening to him? And is a subject likely to change his choices when such considerations underlying them are made fully explicit? Somewhat ironically, it is just because these factors do *not* reflect misinterpretations that they may, after all, help explain some *deliberate* violations.

This is not to say that vagueness, as defined, is typically the sole factor underlying deliberate choices in conflict with the Savage postulates, even in the situations I described, or that such choices reflect mainly a simple aversion to vagueness (though my article may have given those impressions). My own thinking has moved recently toward recognizing the influence of various dimensions of the decision problem under uncertainty that are strongly associated with vagueness but distinct from it; several of Roberts' remarks are highly pertinent and stimulating along these lines.

Nevertheless, as I indicated at the outset, the careful reader of Roberts' catalogue of rationales for violation of the postulates may well find his appreciation of vagueness as a *contributing* factor enhanced rather than diminished. In fact, Roberts' summary remark, "It is hard to see any important role for vagueness . . .," seems to me to make a distinctly odd impression following immediately, as it does, his discussion in Section III. That section ends with a warning which I second, but which also seems cogent in slightly paraphrased form: Conformity to the Savage axioms in spite of vagueness should not be contemplated lightheartedly if the decisions or inferences involved are taken seriously.

THE RAND CORPORATION
SANTA MONICA, CALIFORNIA